

## Building Large Language Models

### Audience Course Building Large Language Models

The course Building Large Language Models is intended for engineers who want to design transformer-based LLMs.

### Prerequisites Course Building Large Language Models

Participants should be comfortable with Python. Prior exposure to PyTorch or a similar Deep Learning framework is helpful.

### Realization Training Building Large Language Models

The training blends concise theory with guided, hands-on labs. Through code-alongs you'll build a mini-GPT, prepare datasets, run pretraining and fine-tuning and deploy models.

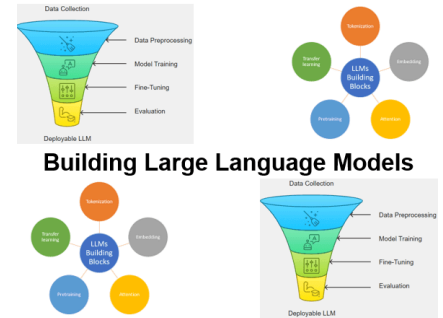
### Building Large Language Models Certificate

After completion, participants receive a certificate of participation for the course Building Large Language Models.

Duration: 4 days

Price: € 3200

[Open Schedule](#)



## Content Building Large Language Models

The course Building Large Language Models from SpiralTrain teaches you how to design, train, fine-tune, and deploy transformer-based LLMs using PyTorch and modern tooling. You will learn to build robust text pipelines and tokenizers, implement a GPT-style model with pretraining, apply instruction tuning and RAG and serve models reliably in production.

### LLM Intro

The course starts by explaining what LLMs are, where they're used, and the lifecycle of building vs. using them. We introduce the Transformer/GPT architecture, how models learn from large datasets, and when to use classic QA versus RAG.

### Working with Text Data

You'll move from raw text to model-ready tensors: tokenization (e.g., BPE), token → ID mapping, special/context tokens, and sliding-window sampling. We cover embeddings and positional encodings while handling unknown words and basic sentence structure.

### Attentions Mechanism

This module demystifies self-attention for long-sequence modeling: queries, keys, values, and causal masking to hide future tokens. We add positional encoding, multi-head attention, and stacked layers to capture dependencies across different parts of the input.

### Pytorch Deep Learning

This module explains PyTorch fundamentals—tensors, core operations, and training loops—with the tooling to measure model quality. We cover feature scaling/normalization (including categorical features), activation and loss functions, and backpropagation.

### Neural Networks

Next the course proceeds with building MLPs and CNNs in PyTorch while choosing appropriate activations and losses and implementing backprop. We touch NLP-specific preprocessing and walk through end-to-end binary and multi-class classification.

### GPT from scratch

Next you will implement a minimal GPT with layer normalization, residual (skip) connections, and attention + feed-forward (GELU) blocks with weight tying. You'll also wire up transformer blocks and generate text to see the model in action.

### Pretraining

Then pretrain the LLM on unlabeled text with next-token prediction, tracking training vs. validation losses. You will explore decoding strategies (e.g., temperature, top-k), control randomness for reproducibility, and save/load PyTorch weights.

### Tuning for Classification

Then the course covers preparing datasets and dataloaders, initializing from pretrained weights, and add a classification head with softmax. Train and evaluate with loss/accuracy, culminating in an LLM-based spam-classification example.

### Fine-Tuning

Finally you will practice supervised instruction tuning: format datasets, batch efficiently, and fine-tune a pretrained LLM. Also evaluate outputs, export responses/checkpoints, and apply parameter efficient methods such as LoRA.

## Modules Building Large Language Models

Module 1 : LLM Intro	Module 2 : Working with Text Data	Module 3 : Attentions Mechanism
What is an LLM? Applications of LLMs Stages of Building LLMs Stages of Using LLMs Transformer Architecture Utilizing Large Datasets GPT Architecture Internals Learn Language Patterns Retrieval Augmented Generation Question and Answer Systems QA versus RAG Building an LLM	Word Embeddings Decoders and Encoders Decoder Only Transformer Tokenizing text Convert Tokens into IDs Special Context Tokens Understand Sentence Structure Byte Pair Encoding Unknown Words Sampling with Sliding Window Creating Token Embeddings Encoding Word Positions	Modeling Long Sequences Capturing Data Dependencies Attention Mechanisms Attending Different Input Parts Using Self-Attention Trainable Weights Hiding Future Words Positional Encoding Causal Attention Masking Weights with Dropout Multihead Attention Stacking Attentions Layers
Module 4 : Pytorch Deep Learning	Module 5 : Neural Networks	Module 6 : GPT from scratch
Deep Learning Intro Overview of PyTorch PyTorch Tensors Tensor Operations Model Evaluation Metrics Feature Scaling Feature Normalization Categorical Features Activation Functions Loss Functions Backpropagation	Neural Networks Intro Building NN with PyTorch Multiple Layers of Arrays Convolutional Neural Networks Activation Functions Loss Functions Backpropagation Natural Language Processing Stopword Removal Binary Classification Multi-class Classification	Coding an LLM Architecture Layer Normalization Normalizing Activations Feed Forward Network GELU Activations Adding Shortcut Connections Connecting Attention Weight Tying Linear Layers in Transformer Block Coding the GPT Model Generating Text
Module 7 : Pretraining	Module 8 : Tuning for Classification	Module 9 : Fine-Tuning
Pretraining on Unlabeled Data Calculating Text Generation Loss Training Losses Validation Set Losses Training an LLM Decoding Strategies Control Randomness Temperature Scaling Saving Model Weights in PyTorch Loading Pretrained Weights	Categories of Fine-Tuning Preparing the Dataset Creating Data Loaders Top-k Sampling Soft-Max Function Initializing with Pretrained Weights Adding Classification Head Classification Loss and Accuracy Fine-tuning on Supervised Data Using LLM as Spam Classifier	Instruction Fine-tuning Supervised Instruction Preparing a Dataset Organizing Training Batches Creating Data Loaders Loading a pretrained LLM Fine-tuning the LLM Extracting and Saving Responses Evaluating Fine-tuned LLM Fine Tuning with LoRA